

WHAT IS CLAIMED IS:

- Sub A^c 7
1. A data processing system comprising:
a source device participating in a multicast group and including:
a first source application instance (AI) producing a first unit of
work stream; and
communication services (CS);
multiple destination devices participating in the multicast group, each
destination device in the multicast group including:
at least one destination AI which consumes units of work; and
CS;
communication services/fabric providing communication between the
source device and the multiple destination devices; and
multiple source and destination resources (SDRs), each SDR
implementing a reliable transport service between the source device and a
corresponding one of the multiple destination devices in the multicast group for
delivery of the first unit of work stream at the corresponding destination device
and guaranteeing strong ordering of the first unit of work stream received at the
corresponding destination device.
 2. The data processing system of claim 1 wherein the CS in the source
device verifies that a predetermined percentage of destination AIs in the
multicast group reliably receives each unit in the first unit of work stream.
 3. The data processing system of claim 2 wherein the predetermined
percentage is 100% of the destination AIs.
 4. The data processing system of claim 2 wherein the predetermined
percentage is less than 100% of the destination AIs.
 5. The data processing system of claim 2 wherein the CS in the source
device includes:

Sub A^c 7

an acknowledgement counter which counts acknowledgements received from the corresponding destination devices in the multicast group indicating that the corresponding destination device has received a unit of work in the first unit of work stream.

5

6. The data processing system of claim 5 wherein the CS in the source device generates a completion event when the acknowledgement counter indicates that the predetermined percentage of destination AIs in the multicast group have acknowledged the unit of work has been received.

10

7. The data processing system of claim 2 wherein the CS in the source device includes:

a bit-mask array which assigns a unique bit for each destination AI in the multicast group and clears each bit as a corresponding acknowledgment is received from the corresponding destination device in the multicast group indicating that the corresponding destination device has received a unit of work in the first unit of work stream.

15

8. The data processing system of claim 7 wherein the CS in the source device generates a completion event when the bit-mask array has the predetermined percentage of bits cleared in the bit-mask array indicating that the predetermined percentage of destination AIs in the multicast group have acknowledged the unit of work has been received.

20

9. The data processing system of claim 1 wherein the CS in the source device replicates the first unit of work stream for transmission to the destination AIs in the multicast group.

25

10. The data processing system of claim 1 wherein the communication services/fabric includes at least one replicater component for replicating the first

30

[illegible]

at least one middleware AI.

13. The data processing system of claim 11 wherein a given AI joins the multicast group by performing a multicast join operation, and the middleware AI or CS determines whether the given AI can join the multicast group, validates access rights, and informs the devices participating in the multicast group of changes in the group.

25

30 16. The data processing system of claim 11 wherein an AI, middleware AI, or CS performs a set attribute operation to set multicast group attributes.

Sub A⁷

17. The data processing system of claim 11 wherein middleware AI performs a remove member operation to remove a given AI from the multicast group without involving the given AI.

5

18. The data processing system of claim 1 wherein an agreed to multicast address is employed to address AIs in the multicast group.

19. The data processing system of claim 18 wherein the CS in each device participating in the multicast group interprets the agreed to multicast address and responds to the agreed to multicast address to perform a reliable multicast operation on behalf of the corresponding destination AI.

20. The data processing system of claim 1 wherein the data processing system performs a reliable multicast operation having substantially the same semantic behavior relative to a given AI as an unreliable multicast operation.

21. The data processing system of claim 1 wherein the multiple SDRs are grouped into multiple SDR groups, wherein each of the multiple SDR groups includes at least one SDR and is assigned a unique priority level for effecting throughput and response time of units of work transmitted by the at least one SDR.

22. The data processing system of claim 1 wherein the source device also functions as a destination device and at least one of the destination devices also functions as a source device.

23. The data processing system of claim 1 wherein each SDR includes:
source SDR resources, at the source device, transmitting the first unit of work stream in a serial unit of work stream having units of work in a defined order over the communication services/fabric; and

Sub A^c 7

destination SDR resources, at the corresponding destination device, receiving the serial unit of work stream, and demultiplexing the serial unit of work stream into units of work provided to the corresponding at least one destination AI.

5

24. The data processing system of claim 23 wherein the destination SDR resources provide a negative acknowledgement (NAK) for a unit of work received ahead of its defined order.

10

25. The data processing system of claim 23 wherein the destination SDR resources drop a unit of work received ahead of its defined order.

15

26. The data processing system of claim 23 wherein the destination SDR resources provide a positive acknowledgement (ACK) for each unit of work which is successfully received and processed by the destination SDR resources.

20

27. The data processing system of claim 23 wherein the destination SDR resources provide a cumulative positive acknowledgement (ACK) for a set of units of work that indicate that all units of work in the set of units of work up to and including a current unit of work have been successfully received and processed by the destination SDR resources.

25

28. The data processing system of claim 18 wherein the destination SDR resources drop a unit of work in response to an indication that the unit of work is a duplicate unit of work.

30

29. A method of processing data comprising the steps of:
producing a first unit of work stream with a first source application instance (AI) at a source device participating in a multicast group;
reproducing the first unit of work stream;

Sub A 7

establishing multiple reliable transport services, each reliable transport service being established between the source device and a corresponding one of multiple destination devices participating in the multicast group, each destination device having at least one destination AI;

- 5 multicasting the reproduced first unit of work stream over a communication services/fabric with the multiple reliable transport services; and guaranteeing strong ordering of the first unit of work stream received at the corresponding destination device.

- 10 30. The method of claim 29 further comprising the step of:
 consuming the first unit of work stream with the at least one destination AI at each of the multiple destination devices participating in the multicast group.

- 15 31. The method of claim 29 further comprising the step of:
 verifying that a predetermined percentage of destination devices in the multicast group reliably has received each unit in the first unit of work stream.

- 20 32. The method of claim 31 wherein the predetermined percentage is 100% of the destination devices.

33. The method of claim 31 wherein the predetermined percentage is less than 100% of the destination devices.

- 25 34. The method of claim 31 further comprising the steps of:
 providing acknowledgments from the corresponding destination devices in the multicast group indicating that the corresponding destination device has received a unit of work in the first unit of work stream; and
 counting acknowledgements received at the source device.

- 30 35. The method of claim 34 further comprising the step of:

Sub A' 7

generating a completion event when the acknowledgement counter indicates that the predetermined percentage of destination devices in the multicast group have acknowledged the unit of work has been received.

5 36. The method of claim 31 further comprising the steps of:
assigning a unique bit for each destination device in the multicast group in a bit-mask array; and
clearing each bit as a corresponding acknowledgment is received from the corresponding destination device in the multicast group indicating that the
10 corresponding destination device has received a unit of work in the first unit of work stream.

37. The method of claim 36 further comprising the step of:
generating a completion event when the bit-mask array has the
15 predetermined percentage of bits cleared in the bit-mask array indicating that the predetermined percentage of destination AIs in the multicast group have acknowledged the unit of work has been received.

38. The method of claim 29 further comprising the step of:
20 maintaining a timing window; and
if the timing window expires without necessary conditions for a completion event occurring, tracking and resolving missing acknowledgments.

39. The method of claim 29 further comprising the steps of:
25 performing a multicast join operation to join a given AI into the multicast group including the steps of:
determining whether the given AI can join the multicast group;
validating access rights; and
informing the devices participating in the multicast group of
30 changes in the group.

Sub A^c 7

40. The method of claim 29 further comprising the steps of:
performing a multicast leave operation to permit a given AI to leave the
multicast group including informing the devices participating in the multicast
group to remove the given AI from the destination list, to complete all in-flight
5 units of work as though the given AI were still present, and to not target the
given AI for units of work not yet launched.
41. The method of claim 29 further comprising the step of:
performing a get attribute operation to query current attributes of the
10 multicast group.
42. The method of claim 29 further comprising the step of:
performing a set attribute operation to set multicast group attributes.
- 15 43. The method of claim 29 further comprising the step of:
performing a remove member operation to remove a given AI from the
multicast group without involving the given AI.
44. The method of claim 29 further comprising the step of:
20 addressing AIs in the multicast group with an agreed to multicast
address.
45. The method of claim 29 further comprising the steps of:
grouping the multiple reliable transport services into multiple reliable
25 transport service groups, wherein each of the multiple reliable transport service
groups includes at least one reliable transport service; and
assigning a unique priority level to each reliable transport service group
for effecting throughput and response time of units of work transmitted by the at
least one reliable transport service.

30

Sub A 7

46. The method of claim 29 wherein the source device also functions as a destination device and at least one of the destination devices also functions as a source device.
- 5 47. The method of claim 29 further comprising the step of:
providing a negative acknowledgement (NAK) for a unit of work
received at a corresponding destination device ahead of a defined order assigned
to the unit of work.
- 10 48. The method of claim 29 further comprising the step of:
dropping a unit of work received at a corresponding destination device
ahead of a defined order assigned to the unit of work.
49. The method of claim 29 further comprising the step of:
15 temporarily storing a unit of work received at a corresponding destination
device ahead of a defined order assigned to the unit of work.
50. The method of claim 29 further comprising the step of:
performing a resynchronization operation to recover a missing
20 intermediate unit of work.
51. The method of claim 29 further comprising the step of:
providing a positive acknowledgement (ACK) for each unit of work
which is successfully received and processed at a corresponding destination
25 device.
52. The method of claim 29 further comprising the step of:
providing a cumulative positive acknowledgement (ACK) for a set of
units of work that indicate that all units of work in the set of units of work up to
30 and including a current unit of work have been successfully received and
processed at a corresponding destination device.

Sub A 7

53. The method of claim 29 further comprising the steps of:
indicating that the unit of work is a duplicate unit of work based on the
unit of work being received at a corresponding destination device behind a
5 defined order assigned to the unit of work; and
dropping the unit of work in response to the indication that the unit of
work is a duplicate unit of work.